DOCUMENT RESUME

ED 147 784 CS 003 827

AUTHOR Thorndyke, Perry W.
TITLE Pattern-Directed Processing of Knowledge from Texts.
INSTITUTION Rand Corp., Santa Monica, Calif.
PUB DATE May 77
NOTE 31p.
AVAILABLE FROM Publications Department, The Rand Corporation, 1700 Main Street, Santa Monica, California 90406 ($3.00)

EDRS PRICE MF-$0.83 HC-$2.06 Plus Postage.
DESCRIPTORS *Artificial Intelligence; Cognitive Processes; *Information Processing; *Learning Processes; Memory; *Reading Comprehension; *Reading Processes; *Reading Research; Recall (Psychological)
IDENTIFIERS *Prose Learning

ABSTRACT

A framework for viewing human text comprehension, memory, and recall is presented that assumes patterns of abstract conceptual relations are used to guide processing. These patterns consist of clusters of knowledge that encode prototypical co-occurrences of situations and events in narrative texts. The patterns are assumed to be a part of a person's world knowledge and can be activated during comprehension to build associations among multiple linguistic propositions in memory according to their higher-order conceptual relations. During text reproduction from memory, these patterns provide retrieval plans for recall and a mechanism for sophisticated "guessing" when retrieval fails. Some data from human text learning tasks are presented as evidence for these higher-order conceptual patterns. Several structural and processing properties of the model are evaluated in light of these data. It is argued that the proposed pattern-directed processing model could be successfully implemented in artificial intelligence systems to provide adaptive error-handling mechanisms such as those observed in human behavior. (Author)

****
* Documents acquired by ERIC include many informal unpublished *
* materials not available from other sources. ERIC makes every effort *
* to obtain the best copy available. Nevertheless, items of marginal *
* reproducibility are often encountered and this affects the quality *
* of the microfiche and hardcopy reproductions ERIC makes available *
* via the ERIC Document Reproduction Service (EDRS). EDRS is not *
* responsible for the quality of the original document. Reproductions *
* supplied by EDRS are the best that can be made from the original. *
****

ED147784

U S DEPARTMENT OF HEALTH,
EDUCATION & WELFARE
NATIONAL INSTITUTE OF
EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

PATTERN-DIRECTED PROCESSING OF KNOWLEDGE FROM TEXTS

Perry W. Thorndyke

May 1977

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

The Rand Corporation

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC) AND USERS OF THE ERIC SYSTEM."

CS003827

P-5806

The Rand Paper Series

Papers are issued by The Rand Corporation as a service to its professional staff. Their purpose is to facilitate the exchange of ideas among those who share the author's research interests; Papers are not reports prepared in fulfillment of Rand's contracts or grants. Views expressed in a Paper are the author's own, and are not necessarily shared by Rand or its research sponsors.

The Rand Corporation
Santa Monica, California 90406

Abstract

A framework for viewing human text comprehension, memory, and recall is presented that assumes patterns of abstract conceptual relations are used to guide processing. These patterns consist of clusters of knowledge that encode prototypical co-occurrences of situations and events in narrative texts. The patterns are assumed to be a part of a person's world knowledge and can be activated during comprehension to build associations among multiple linguistic propositions in memory according to their higher-order conceptual relations. During text reproduction from memory, these patterns provide retrieval plans for recall and a mechanism for sophisticated "guessing" when retrieval fails. Some data from human text learning tasks are presented as evidence for these higher-order conceptual patterns. Several structural and processing properties of the model are evaluated in light of these data. It is argued that the proposed pattern-directed processing model could be successfully implemented in artificial intelligence systems to provide adaptive error-handling mechanisms such as those observed in human behavior.

INTRODUCTION

The comprehension of connected texts has recently emerged as a major area of research in cognitive psychology. Researchers have begun to explore characteristics of people's memory representations of learned prose [5, 7, 10, 11, 13, 14, 17, 26, 27]. Two problems in particular have been addressed in these and other studies of prose memory. The first problem is how to represent the meaning of a text in a way that is more than simply a concatenation of sentences. Clearly, texts have a structure that transcends the individual syntactic structures of the component sentences. People can easily discriminate a well-organized, comprehensible text from a disorganized, confusing one, even when the two texts contain the same knowledge [26]. The second problem is how to account for the fact that people are able to make inferences from texts that require the integration of knowledge across several sentences. At any point during the comprehension of a text, a reader can generate a vast number of plausible inferences. However, people typically generate only those inferences that promote integration of previous and current information, and inferences that establish expectations for subsequent information [25]. Thus, a model of human text memory should explain how and why certain inferences are produced and other equally plausible ones are not.

The purpose of this paper is to argue for certain design considerations in the construction of computer systems that simulate human behavior. Text comprehension provides a

particularly appropriate domain in which to formulate these arguments, since currently humans exhibit considerable success in this task while computers exhibit only limited success. The subsequent discussion is organized as follows. First, a process model is outlined for how people comprehend, store, and retrieve simple narrative texts in memory. The class of texts to be considered is sufficiently constrained to make the task of modeling human memory and performance a reasonable goal, yet it is sufficiently rich to provide a natural processing environment for people. Second, some of the more interesting observations of human performance on memory tasks using these texts are considered and evaluated against predictions of the processing model. These observations will serve to illustrate several desirable properties of successful information processing systems. In particular, it will be argued that these properties can be explicated in terms of pattern-directed recognition and generation mechanisms assumed by the psychological model. Finally, the implications of these observations for the design of artificial intelligence systems to simulate human performance are discussed.

OUTLINE OF A MODEL OF HUMAN TEXT COMPREHENSION

Prototypical Patterns for Story Organization

The model described here was designed to provide a representational scheme for the plot organization of a narrative, that is, the common configurations of situations and events in a

well-formed passage. The class of texts considered by the model includes single-protagonist, single-goal, narrative stories. For example, a simple plot sequence might involve a problem facing the main character, a sequence of episodes in which the main character attempts to solve the problem, and some eventual resolution of the problem. The representational scheme formalizes these relationships as a grammar describing these narrative elements and the dependencies governing their occurrence in combination. The grammar thus provides a framework for encoding, or parsing, the organization of text information according to its narrative structure.

The grammar consists of a set of productions providing rules of narrative syntax and is independent of the linguistic content of any particular story. Each production in the grammar is a pattern of situation and event occurrences that characterizes some critical element in a story, such as the motivation and intention of characters. These patterns are abstract in that they are independent of particular semantic content. Furthermore, the patterns are hierarchically organized: combinations of patterns can organize into higher-order, more inclusive patterns. The successive application of the productions that formalize the structure of the patterns results in a hierarchical representation of a story that includes the semantic content of individual propositions and the organization of these propositions into abstract narrative patterns. Intermediate nodes in the structure represent patterns encoding

narrative relationships among propositions. Terminal nodes represent actual propositions from the story. For example, consider the following simple story:

(1) Judy wanted a new dress for her birthday.

(2) So she went into town shopping for one.

(3) Finally, she found the one she wanted in a small boutique.

(4) She bought the dress and wore it home.

This story may be viewed as a single episode, and the narrative relations among the four events may be represented by the following rule:

(5) Episode --> Goal + Attempt* + Outcome.

The goal, or desired state, can be matched to (1), Judy's desire for a dress. (2) and (3) represent Attempts; that is, events initiated by Judy to attain the goal. (The "*" in (5) indicates that the element may be repeated). The outcome is an event or state that indicates either the success or failure in attaining the goal. Of course, the conceptual inferences required to satisfy the conditions for matching (1)-(4) to the narrative elements are not presented here. For example, the understander must identify (1) as a goal of Judy's, determine that in (2) and (3) "she" refers to Judy and "one" refers to a dress, and that buying the dress in (4) constitutes successful attainment of the goal. The problem of how these inferences are generated is treated elsewhere [4, 20, 21, 23-25]. The

intent here, however, is to show how sequential events can be organized into higher-order patterns to produce a coherent text.

In more complex stories, simple episodes are embedded in a narrative that has a setting, a plot, and some resolution of the plot. The plot is some number of individual episodes comprising attempts to achieve the goal or some subgoal created by the main character. These episodes may be recursively embedded as attempts to resolve subgoals produced during efforts to satisfy higher-order goals. To illustrate the narrative patterns encoding story content, consider the following story:

> (1) Circle Island is located in the middle of the Atlantic Ocean, (2) north of Ronald Island. (3) The main occupations on the island are farming and ranching. (4) Circle Island has good soil, (5) but few rivers and (6) hence a shortage of water. (7) The island is run democratically. (8) All issues are decided by a majority vote of the islanders. (9) The governing body is a senate, (10) whose job is to carry out the will of the majority. (11) Recently, an island scientist discovered a cheap method (12) of converting salt water into fresh water. (13) As a result, the island farmers wanted (14) to build a canal across the island, (15) so that they could use water from the canal (16) to cultivate the island's central region. (17) Therefore, the farmers formed a procanal association (18) and persuaded a few senators (19) to join. (20) The procanal association brought the construction idea to a vote. (21) All the islanders voted. (22) The majority voted in favor of construction. (23) The senate, however, decided that (24) the farmers' proposed canal was ecologically unsound. (25) The senators agreed (26) to build a smaller canal (27) that was 2 feet wide and 1 foot deep. (28) After starting construction on the smaller canal, (29) the islanders discovered that (30) no water would flow into it. (31) Thus the project was abandoned. (32) The farmers were angry (33) because of the failure of the canal project. (34) Civil war appeared inevitable.

By encoding this story according to the narrative grammar, the representation given in Figure 1 is produced. The numbers at

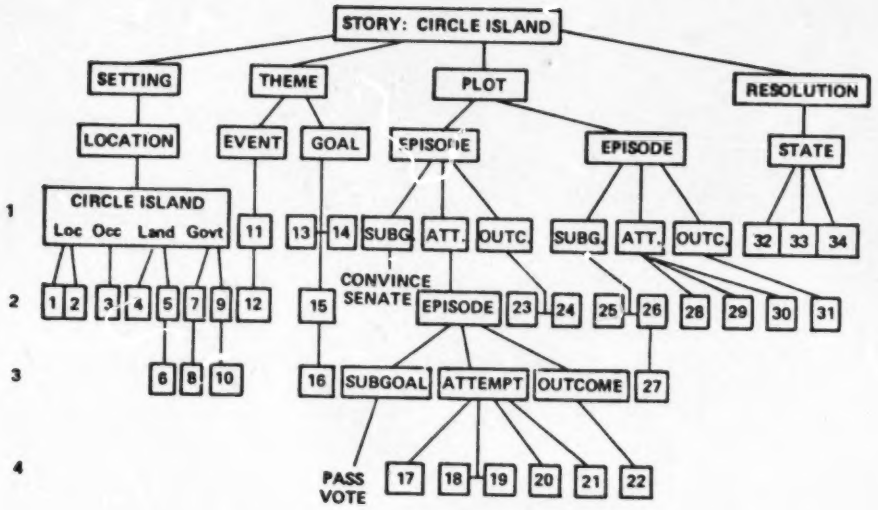

Figure 1. Narrative structure for the Circle Island story.

the terminal nodes of Figure 1 correspond to the numbered propositions given above. The intermediate nodes encode the abstract narrative relations that organize the events of the story into a meaningful narrative according to the constraints specified by the grammar.

This approach to the analysis of text structure presumes a stereotypy in the organization of knowledge in this class of narratives. It further presumes that the conceptual patterns form an integral part of the story representation. While the narrative patterns provide some constraints on the combination of events in a story, the framework can be modified to fit the unique properties of a particular story. Thus stories have invariant properties at some level of abstraction while still varying in the details of particular events and states. This is similar to Bartlett's [2] notion that a story is represented in memory as a general schema and a few specific details. More recently, this general approach to knowledge representation has been embodied in the theories of frames [14], scripts [18, 19], and schemata [16]. However, this theoretical formulation has recurred in experimental psychology for 50 years (cf. [8, 29]).

The Comprehension Process

Comprehension of stories is guided by the patterns of narrative relations. The hierarchical organization of patterns specifies levels of abstraction for matching input to the conceptual structure. The top levels represent the integrating

structural elements of a story, and the lower levels represent specific event and state-recognizing patterns.

Comprehension is viewed as the process of building a representation for the text using the prototypical structural patterns stored in memory. Active processes attempt to match incoming propositions from the story to narrative patterns using the constraints specified by the grammar. For example, to build an episode representation of sentences (1)-(4), matches must be found for a goal, attempts, and an outcome. Obtaining a desired object is frequently a goal, so (1) can be tentatively matched to the goal portion of the episode pattern. Similarly, an effort would be made to match (2)-(4) to the attempt and outcome patterns. The success of these match attempts determines the ease of comprehension. When matches succeed, comprehension is facilitated because relationships among successive events can be inferred from the grammar, and the matched information along with its associated inferences can be represented in familiar well-learned structures in long-term memory [see 25].

Processing proceeds both top-down (goal-directed) and bottom-up (data-driven). At any point during comprehension goal-directed processes actively attempt to match combinations of incoming propositions and instantiated patterns to more general structural patterns by fitting them to expectations provided by the grammar. Simultaneously, data-driven processes attempt to use raw data to fit low-level patterns for states or events to the input. Structural integration is produced by evaluating

these input descriptions in a number of potential contexts and selecting the most reasonable overall interpretation.

Reconstruction from Memory

When a story has been comprehended, it is represented in memory in a hierarchical structure such as the one shown in Figure 1. When asked to recall the story verbatim from memory, a person uses the encoded representation as a retrieval mechanism. His task is to gain access to and retrieve all the terminal nodes from the tree in Figure 1 in correct serial order. Beginning at the highest or most general node in the representation, the subject retrieves information with a depth-first search through the hierarchy. The abstract narrative elements provide intermediate information about the type of knowledge being sought next. In particular, the narrative patterns encode the role of the to-be-recalled knowledge in the story. If the search process encounters a terminal node, that information is output as part of the recalled text.

In most cases, original learning of the passage is less than perfect. This is modeled in the representation as probabilistic strength parameters associated with each link in the hierarchy. Thus, the probability of successful retrieval of information in the hierarchy is a function of the probabilities associated with the path to that information. The deeper in the representational hierarchy a proposition is, then, the lower the probability it can be recalled successfully. Another way to view this prediction is that over time, text-specific knowledge, especially

low-level details of setting and plot, will decay faster than higher-order relational information.

The depth-first retrieval strategy provides a mechanism for recall of the stored information, directed by the narrative patterns encoding the plot structure of the story. When retrieval fails and details cannot be recalled, the subject is able to reconstruct information consistent with the functional role of the unretrieved information, based on the contextual knowledge supplied by the narrative patterns. This allows for a sophisticated guessing strategy, leading to recall errors that are functionally correct in the text as a whole even though the details of the recalled information are incorrect.

SOME OBSERVATIONS OF HUMAN PERFORMANCE

This model of story representation and the associated processes of comprehension and recall makes some interesting predictions for how knowledge is organized in memory, how it is retrieved, and what happens when retrieval attempts fail. In this section a few observations are presented that are consistent with predictions of the model. These observations are given to demonstrate some of the desirable properties of an information-processing system that must deal with imperfect information.

Transition Recall Probabilities

The model assumes that knowledge is organized according to meaningful narrative patterns. These patterns may be viewed as

chunks of information that are conceptually related. If the chunks assumed by the theory are in fact used by people, propositions within one narrative pattern should be more strongly associated than propositions occurring in different patterns.

The data reported here (and later in this section) were obtained from 16 subjects who read and later attempted to recall the story represented in Figure 1. Subjects were presented the story and instructed to recall it in exactly the same order in which it was presented. Thus recall was serial; that is, the input and output sequences were required to be the same. Note that this is the most natural recall order for connected discourse.

Associative strength between propositions in memory was determined by considering the conditional recall probabilities for adjacent propositions in the input sequence. That is, the probability of recall of proposition i is computed given correct recall of proposition i-1. This probability is termed the transition probability between propositions. If subjects are merely storing and recalling propositions independently of each other, the transition probabilities should be the same for all propositions. However, if the proposed model is correct, the transition probability should be low when propositions i-1 and i are in different narrative patterns, and high when they are in the same pattern.

Transition probabilities were first computed between and within the most general narrative constitutents (i.e., Setting,

Theme, Plot, or Resolution). When the transition between two propositions was within one of these constituents (e.g., propositions 30 and 31), the mean conditional recall probability was .44. However, when the transition was between two constituents (e.g., propositions 31 and 32), the probability was .21. When constituents are defined at this most abstract level, the search distance between two adjacent propositions in different constituents is the maximum of all distances between adjacent propositions in the representation.

Now consider the Episode as the basic conceptual constituent. The transition probability for two propositions within an episode was .45. However, the transition probability between adjacent items in two different episodes was .28. Here again a larger network distance between propositions, represented as a between-pattern transition, produced lower transition probabilities than the shorter network distance traversed in within-pattern transitions. These differences are consistent with data from studies of serial recall of hierarchical structures in other verbal learning paradigms (e.g., [9]).

This result suggests that one potential method for associating concepts effectively might be to unify or chunk them; that is, to integrate them into a single higher-order pattern or concept. At the same time, it should be possible to reduce associative strength between concepts by making them constituents of corceptually distinct patterns. This might be an effective storage technique for distinguishing and keeping separate potentially interfering sets

of knowledge.

Intrusion Errors in Recall

When a subject unsuccessfully attempts to recall a portion of the studied story, the failure is often manifested not as a complete omission of information but as recall of information that was not in the original text. These intrusions of information in the recall protocol frequently indicate that the subject has forgotten the details of the story he read but has retained knowledge of the narrative structure of the passage. The evidence for this comes from the fact that the structural role of the intruded information in the narrative is identical to that of the correct information it replaced. That is, the intrusion errors, while incorrect in detail, serve the same function in the overall narrative sequence as the propositions that were actually presented. Consider the following examples:

(6) The senate, however, decided that the farmers' proposed canal was ecologically unsound. The senators agreed to build a smaller canal...

(6') The senate had a vote and rejected the large canal idea and instead proposed a smaller canal...

(7) The islanders discovered that no water would flow into the canal.

(7') It was found that the rivers were dry.

(8) The farmers formed a procanal association and persuaded a few senators to join. The procanal association brought the construction idea to a vote. All the islanders voted. The majority voted in favor of construction.

(8') The farmers asked the senate to support the proposed canal.

The outcome of the first episode of the story is given in (6). (6') gives the recall of the corresponding information from one subject's protocol. Note that while the content of (6) and (6') are entirely different, the effective outcome of the two examples is the same. That is, in both examples the failure of the farmers to achieve their subgoal (convincing the senate to build the canal) is expressed. This knowledge, rather than the details themselves, is critical to the narrative since it provides the justification for the successive episodes.

Similarly, (7) represents part of the attempt to construct the smaller canal taken from the passage. (7') is taken from a subject's protocol. Again the details of the two are quite distinct, yet the failure of the smaller canal is expressed in both cases. In (8') a subject substituted in her recall a single sentence for an entire episode from the text (given in (8)). In this case the episode was embedded in another episode as an attempt to convince the senate to build the canal.

Examples such as these are commonly obtained in recall protocols for stories. In all of these cases subjects appeared to be guided in recall by knowledge of the narrative structure of the story. When memory for detail failed, the subject used the context information provided by the structure in which the information was embedded to generate plausible filler information consistent with the rest of the story. Thus the pattern-directed nature of the recall procedure provided a sort of error-recovery feature that allowed the subject to fill in missing information

and continue his recall of the story.

Hierarchical Knowledge Organization

The story representation in Figure 1 locates the passage content at the terminal nodes in the tree. These propositions appear at different levels in the hierarchy according to the number of productions required to produce the propositions from the top-level node. It may be noted that the higher in the hierarchy a proposition appears, the more global the structural element of the passage it represents. For example, propositions 13 and 14 of Figure 1 refer to the text "The farmers wanted to build a canal across the island." This corresponds to the overall goal of the story, a critical structural element on which subsequent story events depend. Propositions 17-19, on the other hand, refer to the text "The farmers formed a procanal association and persuaded a few senators to join." These propositions represent an attempt to achieve an embedded subgoal and appear lower in the hierarchy. They are, then, less structurally central than propositions 13-14, because their function in the narrative is localized to a low-level pattern. The point here is that the hierarchical organization of knowledge provides an implicit ordering of the importance of information in the narrative structure. Information near the top of the hierarchy represents thematic knowledge of the narrative that is elaborated by details of setting and plot at lower levels in the hierarchy.

This feature of the representation has several implications

for human performance. Since recall proceeds top-down and is a probabilistic function of search distance, recall probability of a proposition should be directly related to its location in the hierarchy. In fact, propositional recall was found to be a monotone decreasing function of hierarchical level. That is, the lower a proposition appeared in the hierarchy, the lower the probability of its recall. Thus the most structurally central information is most reliably recalled, while low-level details are less reliably recalled.

Another implication of the hierarchical organization of knowledge is that any proposition in the structure conceptually subsumes all information directly below it. Lower-order knowledge is a further specification or elaboration of knowledge represented in its higher-order parent. This means, for example, that an entire embedded episode may be omitted in recall, since its role in the narrative can be expressed at a higher level without explicit recall of the episode content. For example, in (8') above, a subject's recall protocol omitted entirely the embedded episode containing propositions 17-22. Instead, the subject inserted a reference to the subgoal of the higher-level episode (winning senate approval), then recalled the outcome of that episode. Although this deletion eliminated an entire episode from the protocol, the resulting transformed story remained a reasonable narrative account.

This knowledge stratification is also a useful framework for understanding subjects' behavior when they are required to

produce concise summaries of the original story. A typical summarization protocol from a subject is given below:

> This passage basically describes the democratic
> attempt of a small island in the Atlantic Ocean [SETTING]
> to build a canal. [GOAL]
> Poor planning [EPISODE 2]
> and internal strife [EPISODE 1]
> finally lead to the project's abandonment. [OUTCOME]
> The people as well as the government are [RESOLUTION]
> enraged
> and civil war seems inevitable between the two [RESOLUTION]
> bodies.

The propositions selected for inclusion in the summary corresponds to the central elements of the narrative structure occurring near the top of the representational hierarchy. That is, the subject included minimal top-level setting information, the goal statement, two word summaries of each of the top-level episodes, the outcome of the second episode, and the story resolution. Details and specific actions from lower levels in the hierarchy have been systematically excluded from the summary. In general, summary content of this story is directly predictable from the representation: only knowledge designated as important (at the top of the hierarchy) will be included in summaries with high probability. In fact, an individual proposition from the top level in the hierarchy containing terminal nodes occurred in

a summary with probability .45; from the second level, .20; and from the third level, .10. Thus the relative importance of story knowledge, reflected in hierarchical location of the knowledge, can be used to select best information to be used in constructing a good summary.

IMPLICATIONS FOR ARTIFICIAL INTELLIGENCE SYSTEMS

In the previous two sections, a number of features of the representation of textual information in memory were presented. It was seen how this model of knowledge representation could be used to predict characteristic behaviors of people operating with that knowledge. It is evident that these behaviors, such as successful comprehension, summarization, error-recovery, and recall exhibit desirable performance properties of an information processing system. Since the model outlined here provides a mechanism for characterizing these behaviors, it seems reasonable that the model's features should be incorporated into the design of computer-based knowledge utilization systems. Some of the features of the human processor that seem particularly well-adapted for use in computer systems are discussed below.

Prototypes as Patterns

The present view of world knowledge is to consider it as a structural framework that guides the acquisition of new knowledge. Thus internal patterns may be considered to be clusters of knowledge that typically co-occur in the world and are often perceived and recognized together. One implication of

this view is that the process of partial-matching input to stored patterns can be used as a basis for inference. For example, if an input sequence has been partially matched to a stored pattern, that pattern will contain residual information that needs to be instantiated to produce a completed match. So an inferential process may be initiated to attempt to match some information in the pattern to some unaccounted for information in the input.

Another type of inference directed by the stored pattern is determined by its association with other patterns. Since patterns combine into higher-order patterns, the recognition of one pattern via a successful match can lead to expectation for matching other patterns in the input sequence. Thus patterns can be used for prediction based on expectancies provided by earlier recognitions. This can effectively order the search for possible interpretations for successive information.

When some of the information that instantiates an instance of a pattern is missing due to its absence from the input or degradation in memory, knowledge of the correct pattern can supply the missing information. This knowledge might be in the form of "default" assignments [12, 15] or in the form of a "guess" that is functionally equivalent to the missing knowledge. This latter case was demonstrated in the recalls of subjects whose protocols contained information that was functionally correct but details that were incorrect. In this case, the fact that the pattern encoding the particular knowledge was well-

learned resulted in better memory for the pattern than for the details instantiating it. This produced a graceful degradation of performance that is characteristic of humans but still an unattained goal for most computer systems. That is, rather than halting completely when information could not be retrieved, an incorrect response was produced that contained the correct functional properties of the missing information.

Hierarchical Knowledge Organization

The concept of organizing knowledge into hierarchies is certainly not a new idea to either computer scientists or psychologists. The use of generalization hierarchies for encoding information leads to convenient storage efficiencies and straightforward deductive mechanisms. In cognitive psychology, numerous models of semantic memory are based on the assumption of hierarchical knowledge representations. The novel concept proposed here is that information about hierarchical location can be used as a metric of information "importance." The knowledge at high levels of the hierarchy encodes the important structural information of the passage as well as the content that corresponds to this structural knowledge. In addition, this high level knowledge conceptually subsumes and integrates other less important clusters of knowledge below it. Since access to the memory representation is assumed to occur at the top of the hierarchy, the most important content is the easiest to retrieve from memory. Less important details, on the other hand, are less readily retrieved but may be produced by searches through the

structure.

This concept of importance ordering might be valuable in situations in which knowledge capacity is limited. A hierarchical representation scheme would permit memory to contain important information in quick access storage while less salient knowledge could be stored remotely and addressed indirectly from conceptual parents. The memory organization scheme could thus be determined by the structural semantics of the knowledge domain.

Patterns as Knowledge Clusters

A fundamental feature of the model presented here is that knowledge units are actually collections of information that combine to produce a single conceptual pattern. For example, an episode is viewed as a pattern of situations, events, and relations among them. This higher-order pattern may be treated as a single unit (as in the deletion of an entire episode from a recall protocol), or it may be decomposed into its components if required.

This view of knowledge representation posits a mechanism for recognizing frequently associated concepts as unitized entities (patterns). The capacity of general patterns to integrate their components provides a basis for strengthening the association among the constituents that instantiate it. Thus, if one desires to associate strongly certain concepts with one another, a technique for accomplishing this might be to unify them under higher-order patterns that require conjunctions of recognized

elements for their instantiation and produce representations reflecting the mutual dependencies.

On the other hand, if one desires to differentiate and maintain similar but distinct concepts, potential interference effects in memory should be minimized. Embedding the concepts in distinct patterns would effectively separate the concepts by distinguishing the contexts in which they were invoked or retrieved. The transition recall probabilities given above for subjects' story recall demonstrated this effect. Associations among items were found to be strong when they were embedded in the same higher-order pattern and weak when embedded in different patterns.

SUMMARY

The preceding discussion has sketched a psychological model for text learning. The model is based on the assumption that learned patterns in memory encode naturally occurring combinations of conceptual elements in narrative sequences. These patterns are used to guide the comprehension, representation, and reproduction of new information encountered through experience. Data have been presented here and elsewhere that are consistent with the predictions of the theory [2, 6, 11, 13, 26]. In addition to the large corpus of confirmatory evidence, substantial theoretical interest has emerged in the concept of prototypical patterns as models of memory [1, 15, 16, 18, 20, 22, 28]. However, the viability of such schematic approaches as models of human memory has not yet been established. On the other hand, while confirmatory evidence in itself is never sufficient to accept a theory, it is the responsibility of researchers with alternative theories to formulate their predictions, demonstrate that their theories fit existing data equally well, and in addition make testable predictions that distinguish their theories from existing ones. To date, no other theories of human story comprehension have emerged satisfying these criteria.

This paper has postulated a schematic representation of knowledge to account for available human data. However, no assumptions have been made that imply a particular representation of the knowledge units. The model described here

could b implemented in software utilizing semantic networks, productions, frames, scripts, actors, etc. From a psychological point of view, the details of any computer implementation of knowledge utilization are unimportant. The important point is that an implemented system should behave in certain ways that are consistent with desirable human behaviors. Ultimately, the only decision criterion for choosing a representation should be how well it works. However, it should be noted that the most successful systems to date have utilized some form of prototypes for pattern recognition, memory encoding, and text generation [3, 20, 24].

References

1. Anderson, R. The notion of schemata and the educational enterprise. In Schooling and the Acquisition of Knowledge, Anderson, R.C., Spiro, R.J., and Montague, W.E. (Eds.), Lawrence Erlbaum Associates, Hillsdale, N.J., 1977.

2. Bartlett, F. Remembering, Cambridge Press, Cambridge, 1932.

3. Bobrow, D. and Winograd, T. An overview of KRL, a knowledge representation language. Cognitive Science, 1, 1977.

4. Charniak, E. Toward a model of children's story comprehension. M.I.T. A.I. Memo TR-266, Massachusetts Institute of Technology, 1972.

5. Crothers, E. Memory structure and the recall of discourse. In Language Comprehension and the Acquisition of Knowledge, Carroll, J. and Freedle, R. (Eds.), V.H. Winston, Washington, 1972.

6. van Dijk, T. and Kintsch, W. Cognitive psychology and discourse: Recalling and summarizing stories. In Trends in Textlinguistics, Dressler, W.U. (Ed.), de Gruyter, New York, 1977.

7. Frederiksen, C. Representing logical and semantic structure of knowledge acquired from discourse. Cognitive Psychology, 7, 1975, 371-458.

8. Head, H. Studies in Neurology, Oxford University Press, New York, 1920.

9. Johnson, N. Organization and the concept of a memory code. In Coding Processes in Human Memory, Melton, A. and Martin, E. (Eds.), V.H. Winston, Washington, 1972.

10. Kintsch, W. Memory representations of text. In Proceedings of the Third Loyola Symposium, Solso, R. (Ed.), Erlbaum, Potomac, Md., 1974.

11. Kintsch, W. Memory for prose. In The Structure of Human Memory, Cofer, C. (Ed.), Freeman, San Francisco, 1975.

12. Kuipers, B. A frame for frames. In Representation and Understanding, Bobrow, D.G. and Collins, A.M. (Eds.), Academic Press, New York, 1975.

13. Mandler, J. and Johnson, N. Rememberance of things parsed: story structure and recall. Cognitive Psychology, 9, 1977, 111-151.

14. Meyer, B.J.F. The Organization of Prose and its Effects on Memory, North-Holland Publishing Co., Amsterdam, 1975.

15. Minsky, M. A framework for representing knowledge. In The Psychology of Computer Vision, Winston, P. (Ed.), McGraw-Hill, New York, 1975.

16. Norman, D. and Bobrow, D. On the role of active memory processes in perception and cognition. In The Structure of Human Memory, Cofer, C. (Ed.), Freeman, San Francisco, 1975.

17. Rumelhart, D.E. Notes on a schema for stories. In Representation and Understanding, Bobrow, D.G. and Collins, A.M. (Eds.), Academic Press, New York, 1975.

18. Rumelhart, D. and Ortony, A. The representation of knowledge in memory. In Schooling and the Acquisition of Knowledge, Anderson, R.C., Spiro, R.J., and Montague, W.E. (Eds.), Lawrence Erlbaum Associates, Hillsdale, N.J., 1977.

19. Schank, R. and Abelson, R. Scripts, plans and knowledge. Proceedings of the Fourth International Joint Conference on Artificial Intelligence, Tbilisi, Georgia, USSR, 1975, 151-157.

20. Schank, R. and Abelson, R. Scripts, Plans, Goals, and Understanding, Lawrence Erlbaum Associates, Hillsdale, N.J., 1977.

21. Schank, R. and Rieger, C. Inference and the computer understanding of natural language. Artificial Intelligence, 5, 1974, 373-412.

22. Schmidt, C. Understanding human action: recognizing the plans and motives of other persons. In Cognition and Social Behavior, Carroll, J. and Payne, J. (Eds.), Lawrence Erlbaum Press, Hillsdale, N.J., 1976.

23. Schmidt, C. and D'Addamio, J. A model of the common-sense theory of intention and personal causation. Proceedings of the Third International Joint Conference on Artificial Intelligence, Stanford, 1973, 465-471.

24. Schmidt, C., Sridharan, N., & Goodson, J. Recognizing plans and summarizing actions. Proceedings of the Conference on Artificial Intelligence and the Simulation of Behavior, Edinburgh, 1976, 291-306.

25. Thorndyke, P. The role of inferences in discourse comprehension. Journal of Verbal Learning and Verbal Behavior, 15, 1976, 437-446.

26. Thorndyke, P. Cognitive structures in comprehension and memory of narrative discourse. Cognitive Psychology, 9, 1977, 77-110.

27. Thorndyke, P. Knowledge transfer in learning from texts. In Proceedings of the NATO International Conference on Cognitive Psychology and Instruction, Plenum Publishers, Amsterdam, 1977.

28. Winograd, T. Frames and the declarative-procedural controversy. In Representation and Understanding, Bobrow, D.G. and Collins, A.M. (Eds.), Academic Press, New York, 1975.

29. Woodworth, R. Experimental Psychology, Holt and Company, New York, 1938.